

CLAIMS

What is claimed is:

1 1. A method of clustering a plurality of network destinations
2 having network addresses being partitioned into groups of network
3 addresses according to an initial grouping, comprising:

4 identifying a plurality of seedpoints from among the network
5 destinations, each seedpoint being an active one of the
6 destinations associated with at least one of the groups of network
7 addresses;

8 topologically clustering the seedpoints into groups of
9 topologically similar seedpoints;

10 performing a measurement from a predetermined location to a
11 seedpoint within each group of seedpoints;

12 clustering the seedpoints into clusters based on the
13 measurements, the clusters being selected in a manner achieving a
14 desired trade-off between the number of clusters and the
15 similarity among the measurements for the seedpoints within each
16 cluster; and

17 generalizing the clusters based on information identifying
18 the network addresses with corresponding seedpoints to which the
19 network addresses are deemed close.

1 2. A method according to claim 1, wherein the seedpoints are
2 identified based on a predetermined desired granularity.

1 3. A method according to claim 2, wherein the predetermined
2 desired granularity is expressed as a number of most significant
3 bits of the destination addresses.

1 4. A method according to claim 3, wherein the destination
2 addresses are 32 bits in length, and the predetermined number of
3 most significant bits is 24.

1 5. A method according to claim 1, wherein the density of the
2 seedpoints in different groups is varied based upon traffic
3 statistics.

1 6. A method according to claim 1, wherein the seedpoints are
2 identified based on information concerning the destinations of
3 data traffic in the network.

1 7. A method according to claim 1, wherein the seedpoints are
2 identified by sending a message to at least one address in each of
3 a complete set of address regions spanning the destination
4 addresses, each region being defined by a corresponding unique
5 pattern of most significant address bits.

1 8. A method according to claim 7, wherein the one address in each
2 of the address regions is one of a set of predetermined addresses
3 within each region to which messages are conditionally sent to
4 identify seedpoints.

1 9. A method according to claim 1, wherein the seedpoints are
2 included in autonomous systems, and further comprising selecting a
3 representative for each cluster of seedpoints, the selecting of a
4 representative including determining whether the representative
5 has the same penultimate hop along a path to the autonomous system
6 as do the seedpoints of the cluster.

1 10. A method according to claim 9, wherein identifying seedpoints
2 includes rejecting those seedpoints for which there is no
3 available representative in the same autonomous system.

1 11. A method according to claim 1, wherein topologically
2 clustering comprises performing traceroute operations to the
3 seedpoints and analyzing the resulting reported routes.

1 12. A method according to claim 1, wherein the measurement
2 performed from the predetermined location to each of the
3 seedpoints is one of multiple measurements performed from the
4 predetermined location to each of the seedpoints.

1 13. A method according to claim 1, wherein the predetermined
2 location is one of multiple predetermined locations from which
3 measurements to the seedpoints are performed.

1 14. A method according to claim 1, wherein performing each
2 measurement comprises sending a time-to-live-limited probe message
3 to a candidate seedpoint.

1 15. A method according to claim 1, wherein performing each
2 measurement comprises sending an echo request message to a
3 candidate seedpoint.

1 16. A method according to claim 1, wherein the measurements are
2 temporal measurements.

1 17. A method according to claim 1, wherein clustering the
2 seedpoints includes ordering the seedpoints according to the
3 measurement.

1 18. A method according to claim 17, further comprising, for each
2 cluster of ordered seedpoints, identifying at least one of the
3 seedpoints whose measurement satisfies a predetermined criterion.

1 19. A method according to claim 18, wherein the predetermined
2 criterion is being closest to a centroid of the measurements of
3 the seedpoints.

1 20. A method according to claim 1, wherein clustering the
2 seedpoints is performed on the basis of autonomous systems in
3 which the seedpoints reside.

1 21. A method according to claim 20, wherein a minimum of one
2 cluster is established per autonomous system.

1 22. A method according to claim 1, wherein clustering the
2 seedpoints is performed on the basis of traffic to the network
3 destinations.

1 23. A method according to claim 1, wherein clustering the
2 seedpoints is performed in at least two passes, a first pass
3 resulting in more clusters than desired, a second pass being based
4 on a subset of the seedpoints taken from larger ones of the
5 clusters resulting from the first pass.

1 24. A method according to claim 1, wherein clustering the
2 seedpoints is based on geographical information about the
3 seedpoints.

1 25. A method according to claim 1, wherein clustering the
2 seedpoints employs a clustering budget of a predetermined number

3 of clusters for each of a predetermined fraction of the total
4 number of seedpoints.

1 26. A method according to claim 1, wherein the clusters are non-
2 overlapping.

1 27. A method according to claim 1, wherein generalizing the
2 clusters results in associating multiple groups of network
3 addresses with each of at least some of the clusters.

1 28. A method according to claim 1, further comprising for each
2 cluster, selecting a representative having a predetermined
3 relationship to the seedpoints of the cluster, and associating the
4 representative with each group of network addresses associated
5 with the cluster.

1 29. A method according to claim 28, wherein the predetermined
2 relationship of each representative to the seedpoints of the
3 associated cluster is a predetermined relationship of the
4 representative to a centroid of the seedpoints.

1 30. A method according to claim 28, wherein the predetermined
2 relationship of each representative to the seedpoints of the
3 associated cluster comprises lying along a network path to a
4 selected one of the seedpoints.

1 31. A method according to claim 28, wherein selecting a
2 representative for the seedpoints of each cluster includes
3 discarding candidate representatives that do not respond to
4 messages.

1 32. A method according to claim 28, wherein selecting a
2 representative for the seedpoints of each cluster includes
3 discarding candidate representatives that respond to messages with
4 high variability.

1 33. A method according to claim 32, wherein thresholds are
2 employed in ascertaining higher-than-acceptable variability, each
3 threshold being associated with a corresponding source of network
4 traffic.

1 34. A method according to claim 28, wherein the representatives
2 are used by an intelligent route controller to select paths for
3 traffic to the destinations, the intelligent route controller
4 being operative to (1) perform periodic measurements to each of
5 the representatives via different connections of the intelligent
6 route controller, and (2) on the basis of the periodic
7 measurements to the representatives, conditionally modify which of
8 the connections is used for traffic sent to the network
9 destinations.

1 35. A method according to claim 1, wherein the initial grouping of
2 network addresses is established by a set of address prefixes.

1 36. A method according to claim 35, wherein the address prefixes
2 are also employed to establish closeness in the generalizing step.

1 37. A method according to claim 35, wherein the address prefixes
2 reside in a routing table.

1 38. A method of clustering a plurality of network destinations
2 having addresses spanned by a set of address prefixes, comprising:

3 identifying a plurality of seedpoints from among the network
4 destinations, each seedpoint being an active one of the
5 destinations associated with a corresponding at least one of the
6 address prefixes;

7 topologically clustering the seedpoints into groups of
8 topologically similar seedpoints;

9 performing a measurement from a predetermined location to a
10 seedpoint within each group of seedpoints;

11 clustering the seedpoints into clusters based on the
12 measurements, the clusters being selected in a manner achieving a
13 desired trade-off between the number of clusters and the
14 similarity among the measurements for the seedpoints within each
15 cluster;

16 generalizing the clusters based on the address prefixes, the
17 generalizing including conditionally modifying the set of address
18 prefixes such that each address prefix in the conditionally
19 modified set of address prefixes is associated with a
20 corresponding single one of the clusters.

1 39. A method according to claim 38, wherein the density of the
2 seedpoints in different address prefixes is varied based upon
3 traffic statistics.

1 40. A method according to claim 38, wherein generalizing the
2 clusters includes associating each seedpoint with the longest one
3 of those address prefixes matching the seedpoint.

1 41. A method according to claim 38, wherein generalizing the
2 clusters results in associating multiple address prefixes with
3 each of at least some of the clusters.

1 42. A method according to claim 38, wherein conditionally
2 modifying the set of address prefixes comprises recursively
3 splitting each address prefix that matches seedpoints from
4 multiple clusters until each resulting address prefix matches
5 seedpoints from only one cluster.

1 43. A method according to claim 38, wherein conditionally
2 modifying the set of address prefixes comprises recursively
3 merging address prefixes having greater granularity than the
4 address prefixes in the set of address prefixes until any further
5 merging would result in associating at least one address prefix
6 with seedpoints of multiple clusters.

1 44. A method according to claim 38, further comprising for each
2 cluster, selecting a representative having a predetermined
3 relationship to the seedpoints of the cluster, and associating the
4 representative with each address prefix associated with the
5 cluster in the conditionally modified set of address prefixes.

1 45. A method according to claim 44, wherein the predetermined
2 relationship of each representative to the seedpoints of the
3 associated cluster is a predetermined relationship of the
4 representative to a centroid of the seedpoints.

1 46. A method according to claim 44, wherein the predetermined
2 relationship of each representative to the seedpoints of the
3 associated cluster comprises lying along a network path to a
4 selected one of the seedpoints.

1 47. A method according to claim 44, wherein selecting a
2 representative for the seedpoints of each cluster includes

3 discarding candidate representatives that do no respond to
4 messages.

1 48. A method according to claim 44, wherein selecting a
2 representative for the seedpoints of each cluster includes
3 discarding candidate representatives that respond to messages with
4 high variability.

1 49. A method according to claim 48, wherein thresholds are
2 employed in ascertaining higher-than-acceptable variability, each
3 threshold being associated with a corresponding source of network
4 traffic.

1 50. A method according to claim 44, wherein the representatives
2 are used by an intelligent route controller to select paths for
3 traffic to the destinations, the intelligent route controller
4 being operative to (1) perform periodic measurements to each of
5 the representatives via different connections of the intelligent
6 route controller, and (2) on the basis of the periodic
7 measurements to the representatives, conditionally modify which
8 of the connections is used for traffic sent to the network
9 destinations.